

BizSim

The World of Business—in a Box

JOHN L. CASTI

THE COMPUTER AS A LABORATORY

The central process distinguishing science from its competitors—religion, music, literature, mysticism—in the reality-generation business is the so-called *scientific method*. An integral part of this method by which we arrive at scientific “truth,” is the ability to do controlled, repeatable laboratory experiments by which hypotheses about the phenomenon under investigation can be tested. It is just such experiments that on a good day lead to the theories and paradigms constituting today’s “scientific” world view. And, more than anything else, it is the inability to perform experiments of this type that separate the natural sciences from the worlds of social and behavioral phenomena. In the latter, we have no way of doing the experiments necessary to create a bona fide scientific theory of processes like stock market dynamics, road-traffic flow, and organizational restructuring.

In an earlier, less discerning era, it was often claimed that the realm of human social behavior was beyond the bounds of scientific analysis, simply because human beings are “complex”, “unpredictable”, “display free will”, “act randomly”, and so on and so forth. It’s hard to believe that any modern system theorist would do anything but laugh at such childish and naïve attitudes to the creation of workable and worthwhile *scientific* theories of social and behavioral phenomena. The major barrier to bringing the social beneath the umbrella of science is not the non-explanations just given in quotes, but the fact that until now we have had no way to test hypotheses and, therefore, make use of the scientific method in the creation of theories of social behavior. Now we do. And that laboratory in which we do our experiments is the digital computer. Let me illustrate with an example from the world of finance.

BOOMS AND BUSTS, BUBBLES AND CRASHES

In the fall of 1987, W. Brian Arthur, an economist from Stanford, and John Holland, a computer scientist from the University of Michigan, were sharing a house in Santa Fe while both were visiting the Santa Fe Institute. During endless hours of evening conversations over numerous beers, Arthur and Holland hit upon the idea of creating an artificial stock market inside a computer, one that could be used to answer a number of questions that people in finance had wondered and worried about for decades. Among those questions were:

- Does the average price of a stock settle down to its so-called *fundamental value*—the value determined by the discounted stream of dividends that one can expect to receive by holding the stock indefinitely?
- Is it possible to concoct technical trading schemes that systematically turn a profit greater than a simple buy-and-hold strategy?
- Does the market eventually settle into a fixed pattern of buying and selling? In other words, does it reach “stationarity”?
- Alternately, does a rich “ecology” of trading rules and price movements emerge in the market?

Arthur and Holland knew that the conventional wisdom of finance argued that today’s price of a stock is simply the discounted *expectation* of tomorrow’s price plus the dividend, given the information available about the stock today. This theoretical price-setting procedure is based on the assumption that there is an objective way to use today’s information to form this expectation. But the information available typically consists of past prices, trading volumes, economic indicators, and the like. So there may be many perfectly defensible ways based on many different assumptions to statistically process this information in order to forecast tomorrow’s price. For example, we could say that tomorrow’s price will equal today’s price. Or we might predict that the new price will be today’s price divided by the dividend rate. And so on and so forth.

The simple observation that there is no single, best way to process information led Arthur and Holland to the not-very-surprising conclusion that deductive methods for forecasting prices are, at best, an academic fiction. As soon as you admit the possibility that not all traders in the market arrive at their forecasts in the same way, the deductive approach of classical finance theory, which relies upon following a *fixed* set of rules to determine tomorrow's price, begins to break down. So a trader must make assumptions about how other investors form expectations and how they behave. He or she must try to psyche out the market. But this leads to a world of *subjective* beliefs and to beliefs about those beliefs. In short, it leads to a world of induction in which we generalize rules from specific observations rather than one of deduction.

In order to address these kinds of questions, Arthur, Holland and their colleagues constructed an electronic stock market, in which they could manipulating trader's strategies, market parameters, and all the other things that cannot be done with real exchanges. The traders in this market are assumed to each summarize recent market activity by a collection of descriptors, which involve verbal characterization like "the price has gone up every day for the past week," or "the price is higher than the fundamental value," or "the trading volume is high." Let us label these descriptors A, B, C , and so on. In terms of the descriptors, the traders decide whether to buy or sell by rules of the form: "If the market fulfills conditions A, B , and C , then buy, but if conditions D, G, S , and K are fulfilled, then hold." Each trader has a collection of such rules, and acts in accordance with only one rule at any given time period. This rule is the one that the trader views as his or her currently most accurate rule.

As buying and selling goes on in the market, the traders can reevaluate their different rules by assigning higher probability of triggering a given rule that has proved profitable in the past, and/or by recombining successful rules to form new ones that can then be tested in the market. This latter process is carried out by use of what is called a genetic algorithm, which mimics the way nature combines the genetic pattern of males and females of a species to form a new genome that is a combination of those from the two parents.

A run of such a simulation involves initially assigning sets of predictors to the traders at random, and then beginning the simulation with a particular history of stock prices,

interest rates, and dividends. The traders then randomly choose one of their rules and use it to start the buying-and-selling process. As a result of what happens on the first round of trading, the traders modify their estimate of the goodness of their collection of rules, generate new rules (possibly), and then choose the best rule for the next round of trading. And so the process goes, period after period, buying, selling, placing money in bonds, modifying and generating rules, estimating how good the rules are, and, in general, acting in the same way that traders act in real financial markets.

A typical frozen moment in this artificial market is displayed in Figure 1. Moving clockwise from the upper left, the first window shows the time history of the stock price and dividend, where the current price of the stock is the black line and the top of the grey region is the current fundamental value. The region where the black line is much greater than the height of the grey region represents a price bubble, whereas the market has crashed in the region where the black line sinks far below the grey. The upper right window is the current relative wealth of the various traders, and the lower right window displays their current level of stock holdings. The lower left window shows the trading volume, where grey is the number of shares offered for sale and black is the number of shares that traders have offered to buy. The total number of trades possible is then the smaller of these two quantities, because for every share purchased there must be one share available for sale.

After many time periods of trading and modification of the traders' decision rules, what emerges is a kind of ecology of predictors, with different traders employing different rules to make their decisions. Furthermore, it is observed that the stock price always settles down to a random fluctuation about its fundamental value. However, within these fluctuations a very rich behavior is seen: price bubbles and crashes, market moods, over-reactions to price movements, and all the other things associated with speculative markets in the real world.

The agents in the stockmarket simulation are individual traders. A quite different type of business simulation emerges when we want to look at an entire industry, in which case the agents become the individual firms constituting that industry. The world's catastrophe insurance industry served as the focus for just such a simulation exercise called *Insurance*

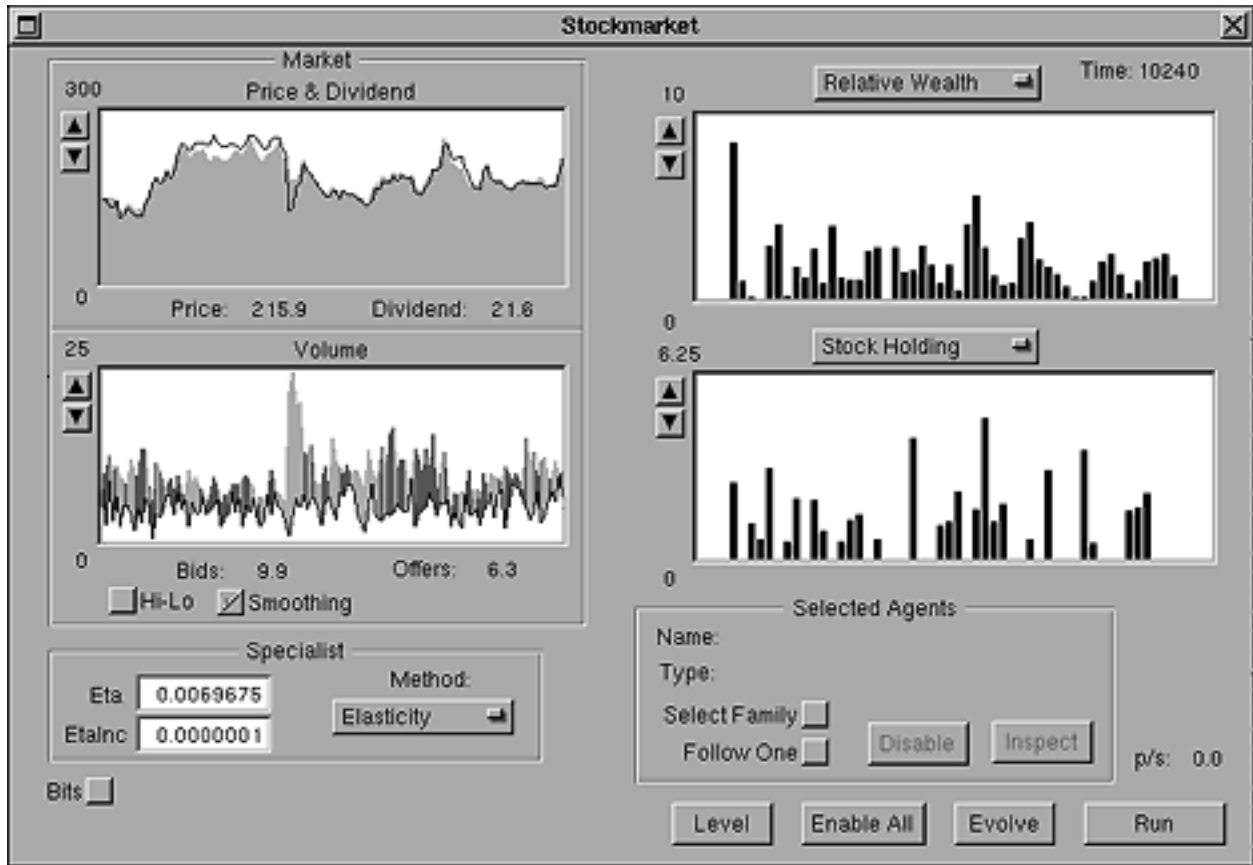


Figure 1. A frozen moment in the surrogate stock market.

World, carried out by the author and colleagues at the Santa Fe Institute and Intelligize, Inc. over the past couple of years.

INSURANCE WORLD

As a crude, first-cut, the insurance industry can be regarded as an interplay among three components: *firms*, which offer insurance, *clients*, who buy it, and *events*, which determine the outcomes of the “bets” that have been placed between the insurers and their clients. In Insurance World, the agents consist of primary casualty insurers and the reinsurers, the firms that insure the insurers, so to speak. The events are natural hazards, such as hurricanes and earthquakes, as well as various external factors like government regulators and the global capital markets.

Insurance World is a laboratory for studying questions of the following sort:

- *Optimal Uncertainty:* While insurers and reinsurers talk about getting a better handle on uncertainty so as to more accurately assess their risk and more profitably price their product, it's self-evident that perfect foreknowledge of natural hazards would spell the end of the insurance industry. On the other hand, complete ignorance of hazards is also pretty bad news, since it means there is no way to weight the bets the firms make and price their product. This simple observation suggests that there is some optimal level of uncertainty at which the insurance—but perhaps not their clients—can operate in the most profitable and efficient fashion. What is that level? Does it vary across firms? Does it vary between reinsurers, primary insurers, and/or end consumers?

- *Industry Structure:* In terms of the standard metaphors used to characterize organizations—a machine, a brain, an organism, a culture, a political system, a psychic prison—which type(s) most accurately represents the insurance industry? And how is this picture of the organization shaped by the specific “routines” used by the decisionmakers in the various components making up the organization?

The simulator calls for the management of each firm to set a variety of parameters having to do with their desired market share in certain regions for different types of hazards and level of risk they want to take on, as well as to provide a picture of the external economic climate (interest rates, likelihood of hurricanes/earthquakes, inflation rates and so forth). The simulation then runs for 10 years in steps of one quarter, at which time a variety of outputs can be examined. For instance, Figure 2 shows the market share for Gulf Coast hurricane insurance of the five primary insurers in this toy world, under the assumption that the initial market shares were *almost* identical—but not quite. In this experiment, firm 2 has a little larger initial market share than any of the other firms, a differential advantage that it then uses to squeeze out *all* the other firms at the end of the ten-year period. This is due to the “brand effect,” in which buyers tend to purchase insurance from companies that they know about.

As a final example of what simulation and business have to say to each other, consider the movement of shoppers in a typical supermarket. This world is dubbed *SimStore* by Ugur Bilge of SimWorld, Ltd. and Mark Venables at J. Sainsbury in London, who collaborated with the author on its creation.




Figure 2. Market share distribution for five primary insurers.

SIMSTORE

The starting point for SimStore is a real supermarket in the Sainsbury chain, one located in the London region of South Ruislip. The agents are individual shoppers who frequent this store. These electronic shoppers are dropped into the store, and then make their way to the various locations in the store by rules such as “wherever you are now, go to the location of the nearest item on your shopping list,” so as to gather all the items they want to purchase.

As an example of one of the types of outputs generated by SimStore, customer checkout data are used to calculate customer densities at each location. Color codes are with descending order: blue, red, purple, orange, pink, green, cyan, grey and nothing. Using the Manhattan metric pattern of movement, in which a customer can only move along the aisles of the store, all locations above 30 percent of customer densities have been linked to form a most popular customer path. Once this path is formed a genetic algorithm will minimize (or maximize!) the length of the overall shopping path.

In the same store, this time each individual customer path has been internally calculated using the simple “nearest neighbor” rule noted above. All customer paths have been summed for each aisle, in order to calculate the customer path densities. These densities are displayed in Figure 3 as a relative density map using the same color code just mentioned.

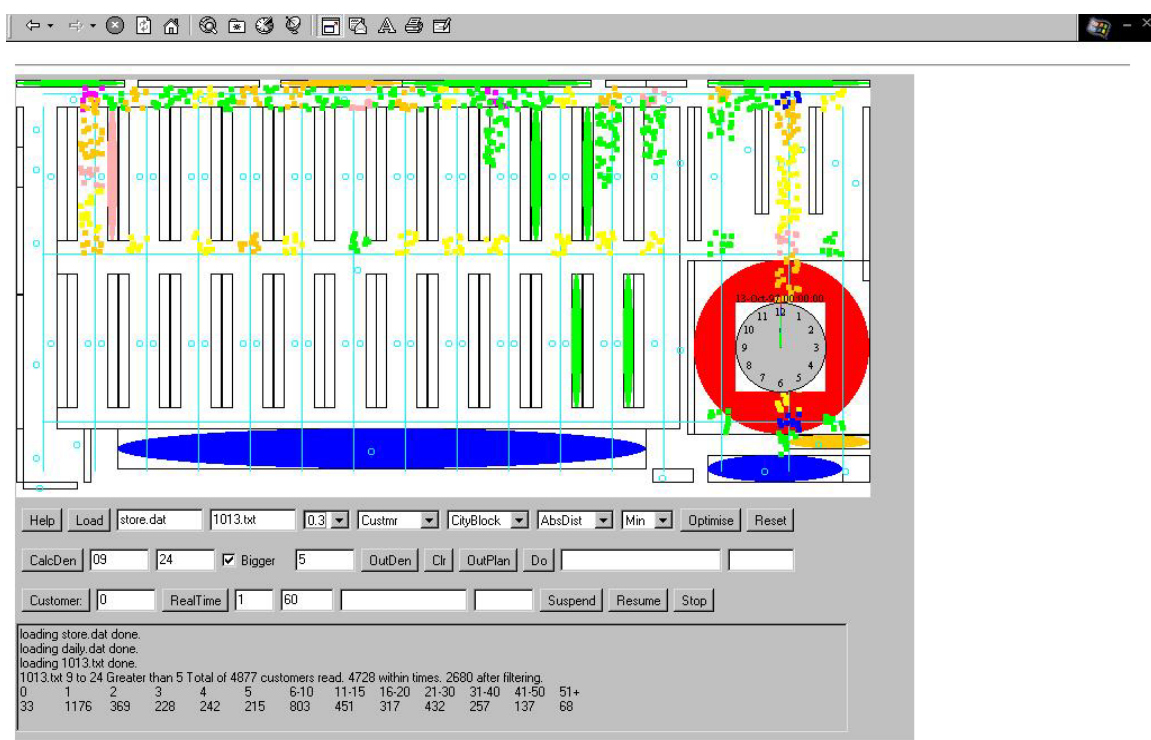


Figure 3. Customer densities along each aisle in the simulated store.

SIMULATION IS GOOD FOR BUSINESS

Large-scale, agent-based simulations of the type discussed here are in their infancy. But even the preliminary exercises outlined here show the promise of using modern computing technology to provide the basis for doing experiments that have never been possible before. Even better, these experiments are exactly the sort called for by the scientific method—controlled and repeatable—so that for the first time in history we have the opportunity to actually create a *science* of human affairs. If I were placing bets on the matter, I’d guess

that the world of business and commerce will lead the charge into this new science that will form during the 21st century.

References

Casti, J. *Would-Be Worlds*. New York: Wiley, 1997.